arXiv:1411.0602v1 [cs.LG] 3 Nov 2014

Factorbird - a Parameter Server Approach to Distributed Matrix Factorization

Sebastian Schelter*
Technische Universität Berlin
@sscdotopen

Venu Satuluri
Twitter
@vsatuluri

Reza Bosagh Zadeh
Stanford University
@reza_zadeh

Abstract

We present 'Factorbird', a prototype of a parameter server approach for factorizing large matrices with Stochastic Gradient Descent-based algorithms. We designed Factorbird to meet the following desiderata: (a) scalability to tall and wide matrices with dozens of billions of non-zeros, (b) extensibility to different kinds of models and loss functions as long as they can be optimized using Stochastic Gradient Descent (SGD), and (c) adaptability to both batch and streaming scenarios. Factorbird uses a parameter server in order to scale to models that exceed the memory of an individual machine, and employs lock-free Hogwild!-style learning with a special partitioning scheme to drastically reduce conflicting updates. We also discuss other aspects of the design of our system such as how to efficiently grid search for hyperparameters at scale. We present experiments of Factorbird on a matrix built from a subset of Twitter's interaction graph, consisting of more than 38 billion non-zeros and about 200 million rows and columns, which is to the best of our knowledge the largest matrix on which factorization results have been reported in the literature.

1 Introduction

In recent years, there is a growing trend to apply machine learning (ML) to massive and complex datasets [1]. For many applications, this leads to growing model sizes. A prime example are recommender systems where the size of the model is usually proportional to the number of users in the dataset [2, 9]. In cases with hundreds of millions of users this model exceeds the memory of an individual machine and distributed approaches to model training that leverage multiple machines become necessary. This leads to a set of challenges, e.g. how to partition the model among the participating machines and how to correctly execute learning algorithms in such a setting.
In this work, we describe 'Factorbird', a prototypical system that leverages a parameter server architecture [17] for learning large matrix factorization models for recommendation mining. After a short introduction to matrix factorization for recommender systems (Sec. 2), we describe the main challenges of our system, namely how to partition the model among the machines in the cluster and how to run Stochastic Gradient Descent (SGD) in parallel. We discuss the design decisions taken in Factorbird to overcome these challenges (Sec. 3). First, we partition one of the matrices to learn over dedicated machines in the cluster and co-partition the other one with the input data to localize a large number of updates and drastically reduce update conflicts and network traffic. Next, we apply a lock-less Hogwild!-style [14] execution scheme to efficiently run SGD in parallel. After giving insights into our software design and memory-efficient datastructures (Sec. 4), we describe techniques to assess model quality in Factorbird (Sec. 5). Our approach here is to grid search over a large number of hyperparameter combinations in a single training run. Finally, we present a set of experiments on user interaction data from twitter (Sec. 6). We run a scale-out experiment on a matrix built from a subset of Twitter's interaction graph [20], with more than 38 billion non-zeros

*work done while at Twitter

and about 200 million rows and columns. To the best of our knowledge, this is largest dataset on which matrix factorization experiments have been published so far.

2 Background: Latent Factor Models

We build upon *latent factor models*, which leverage a low-rank matrix factorization of interaction data [2] to characterize users and items by vectors of factors inferred from interaction patterns. These methods have been among the top-scoring contributions to the Netflix prize [3]. They factor a sparse partially-observed $m \times n$ matrix M, representing the interactions of m users with n items, into the product of two rank k factor matrices U and V, such that their product UV approximates the observed parts of M and generalizes well to unobserved parts of M (c.f. Figure 1). A user i is associated to a factor vector $u_i \in \mathbb{R}^k$ (a row of the $m \times k$ matrix U), while an item j is associated to a factor vector $v_j \in \mathbb{R}^k$ (a column of the $n \times k$ matrix V). The factorization maps users and items onto joint latent factor space of low dimensionality k, such that $u_i^T v_j$ estimates the strength of interaction between user i and item j.

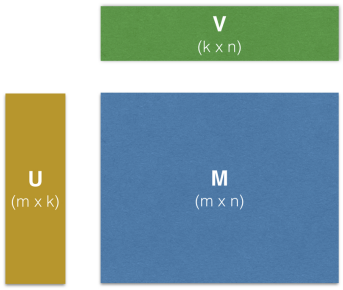

Figure 1: Latent factor models as matrix factorization problem.

The standard approach to learning a latent factor model for recommender systems is to minimize the regularized squared error over the predictions to the observed parts of the matrix [2]:

$$\min_{U,V} \sum_{(i,j) \in M} (m_{ij} - u_i^T v_j)^2 + \lambda \left(\|u_i\|^2 + \|v_j\|^2 \right)$$

This approach is closely related to the Singular Value Decomposition (SVD) of the matrix, which gives an optimal rank k approximation (w.r.t. to squared error) of the original matrix using the top-k singular values and singular vectors. The key difference here is that the SVD is undefined when there are missing entries.

We adopt a more sophisticated approach outlined in [2]. First, we introduce a global bias term g, which captures the average interaction strength between a user and an item in the dataset. Next, we introduce a user-specific bias term b_i^U for user i and an item-specific bias term bias b_j^V for item j. These biases model how strongly the interactions of certain users and items tend to deviate from the global average. We substitute the dot product $u_i^T v_j$ for predicting the strength of the interaction between a user i and an item j with a prediction function $p(i,j) = g + b_i^U + b_j^V + u_i^T v_j$ that takes the bias terms into account. Furthermore, we introduce a function $a(i,j)$ for determining the strength of the interaction between user i and item j. This allows us to transform the observed entries m_{ij} of M. Finally, we add a function $w(i,j)$ for weighting the prediction error for interaction between user i and item j. This function becomes useful when M consists of data from various sources with differing confidence in the observations. The loss function that we minimize for our latent factor model in Factorbird is the following:

$$\min_{g,b^U,b^V,U,V} \frac{1}{2} \left(\sum_{i,j \in M} w(i,j)(p(i,j) - a(i,j))^2 \right) + \frac{\lambda}{2} \left(g^2 + \|b^U\|^2 + \|b^V\|^2 + \|U\|_F^2 + \|V\|_F^2 \right)$$

We adopt a graph-specific terminology for the rest of the paper as well as for the APIs of our system, as we think about the majority of Twitter's datasets in terms of large networks. Therefore,

we assume that M represents a network (e.g. the network of user followings in Twitter), i and j reference vertices in this network (e.g. two users in this network) and the edges of the network correspond to observed entries of M, meaning that $a(i, j)$ depicts the weight of a directed edge between a vertex i and a vertex j in this network. Furthermore, we assume that the bias vectors b^U and b^V are stored in U and V (e.g., as first column and first row), to simplify notation.

Factorization via Stochastic Gradient Descent (SGD). There are various popular techniques to compute a matrix factorization. Stochastic Gradient Descent (SGD) [2] randomly loops through all observed interactions, computes the error of the prediction for each interaction and modifies the model parameters in the opposite direction of the gradient. Another technique is Alternating Least Squares (ALS) [8, 16, 9], which repeatedly keeps one of the unknown matrices fixed, so that the other one can be optimally re-computed. We chose SGD as optimization technique, as it is simple, provides fast convergence and is easy to adapt to different models and loss functions. Algorithm 1 shows the individual steps to conduct when learning the matrix factorization with SGD. First, we randomly initialize the factor matrices U and V (c.f. line 1). Next, we randomly pick an edge (i, j) from the graph and compute the weighted error e_{ij} of the prediction $p(i, j)$ against the actual edge strength $a(i, j)$ (c.f. lines 3 & 4). Next, we update the global bias term as well as the bias terms for i and j proportional to the prediction error e_{ij}, the learning rate η and the regularization constant λ (c.f. lines 5 to 7). We weight the regularization updates according the out-degree n_i of vertex i (the number of observed entries in the i-th row of M) and the in-degree n_j of j (the number of observed entries in the j-th column of M). We update the factor vectors u_i and v_j analogously (c.f. lines 8 & 9). The whole process is repeated until convergence.

Algorithm 1: Matrix Factorization using SGD.

1 randomly initialize U and V

2

3 **while** *not converged* **do**

4 randomly pick edge (i, j)

5

// compute weighted prediction error

6 $e_{ij} \leftarrow w(i,j)(a(i,j) - p(i,j))$

7

// update biases

8 $g \leftarrow g - \eta\,(e_{ij} + \lambda g)$

9 $b_i^U \leftarrow b_i^U - \eta\left(e_{ij} + \frac{\lambda}{n_i} b_i^U\right)$

10 $b_j^V \leftarrow b_j^V - \eta\left(e_{ij} + \frac{\lambda}{n_j} b_j^V\right)$

11

// update factors

12 $u_i \leftarrow u_i - \eta\left(e_{ij}\; v_j + \frac{\lambda}{n_i} u_i\right)$

13 $v_j \leftarrow v_j - \eta\left(e_{ij}\; u_i + \frac{\lambda}{n_j} v_j\right)$

3 System Design

Having introduced the conceptual background of the models that we wish to learn, we proceed with describing the three main design goals of Factorbird.

First, Factorbird must handle factorizations of twitter-scale graphs with hundreds of millions of vertices and dozens of billions of edges. Scalability to datasets of this size is more important than high performance on small datasets commonly used in research (such as the Netflix [3] dataset). The matrices representing these graphs are either square (e.g. user to user followings) or 'tall-and-wide' for bipartite graphs (e.g. user to tweets). Second, the system has to be highly usable and adaptable for data scientists without making a systems background a necessity. No systems programming should be required to try out a variation of our model or a different loss function. Instead, the system should offer interfaces that abstract from the distributed execution model and

are intuitive to implement given an ML background. Third, the system design shall be simple to keep the maintenance and debugging effort low. Additionally, the system design should be easily extendable into a streaming system in the future, where a previously learned matrix factorization is updated online given a stream of new observations. This is the main reason why decided against using ALS, which is much harder to adapt to a streaming scenario than SGD.

Challenges. Running SGD on datasets of such scale inevitably leads to a set of challenges that have to be overcome when designing a scalable system. The main focus of this paper is to present the prototype of a system that elegantly solves these challenges.

(1) The resulting factor matrices for a huge network quickly become larger than the memory available on an individual commodity machine [22, 23]. For example, U and V with $k = 100$ and a single precision factor representation for a graph with 250 million vertices already have a combined size of about 200 GB. This estimation does not even take operating system buffers and caches, object references and required statistics like degree distribution of the graph into account, which also compete for memory.

(2) Due to the sheer number of observations in large datasets, we aim to leverage multiple cores and machines to learn our desired matrix factorization. Unfortunately, SGD is an inherently sequential algorithm. It randomly picks an observation and updates the model parameters before proceeding to the next observation (c.f. Algorithm 1). When we run SGD in parallel on multiple cores, there is a chance that we concurrently try to update the same u_i or v_j, which results in conflicting writes and lost updates.

In order to overcome the first challenge, we decided for a distributed architecture that allows us to partition the large model (the factor matrices) over several machines. We adapt a 'Parameter Server' architecture [17]. As illustrated in Figure 2, we partition the factor matrices over a set of machines, to which we refer to as *parameter machines*. At the same time, we partition the graph (our input data) over a set of so-called *learner machines*. Each learner machine runs multi-threaded SGD on its portions of the input data. For every observation to process, the learner machine has to fetch the corresponding factor vectors from the parameter machines, update them and write them back over the network.

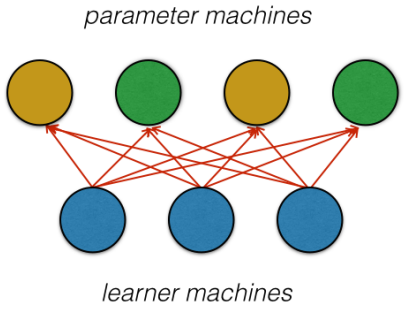

Figure 2: Parameter Server architecture.

This architecture inevitably leads to the second challege, the question of how to handle concurrent, possibly conflicting updates to the factor vectors. When two learner machines fetch, update and write back the same factor vector concurrently, one such update will be overridden.
In the special case of matrix factorization, approaches to parallelizing SGD have been proposed that leverage a carefully chosen partitioning of the input data to avoid conflicting updates [7, 15]. As these approaches require complex data movement and synchronization patterns and are at the same time hard to adapt to a streaming scenario, we decided for an alternative approach that is simpler to implement in a distributed setting. Instead of taking complex actions to prevent conflicting updates, Factorbird builds upon a recently proposed parallelization scheme for SGD-based learning called Hogwild! [14]. This work states that parallel SGD can be implemented *without any locking* if most updates only modify small parts of the model. The authors explicitly name latent factor models for matrix factorization as one such case.
The special nature of the matrix factorization problem allows us for a further optimization in our

system design, which reduces the required network traffic and at the same time greatly lowers the probability for conflicting overwrites. We can reduce the communication cost by 50% through intelligent partitioning, as follows. If we partition M by either rows or columns over the learning machines, than the updates to either U or V become local, when we co-partition one of the factor matrices (U in the case of partitioning by rows, V in the case of partitioning by columns) with M on the learner machines.
In the light of minimizing update conflict potential, we decide to co-locate V on the learner machines and keep U in the parameter machines (c.f. Figure 3). We choose this scheme for the following reasons: In case of the follower graph, the number of updates to a factor vector u_i in U is equal to the out-degree of the corresponding vertex i, while the number of updates to a factor vector v_j in V is equal to the in-degree of the corresponding vertex j. As the in-degree distribution of the follower graph has much higher skew than the out-degree distribution [21], we choose to localize the updates to V, which gives us a higher reduction in conflict potential than localizing U. Other graphs in twitter have similar skew in the degree distribution.

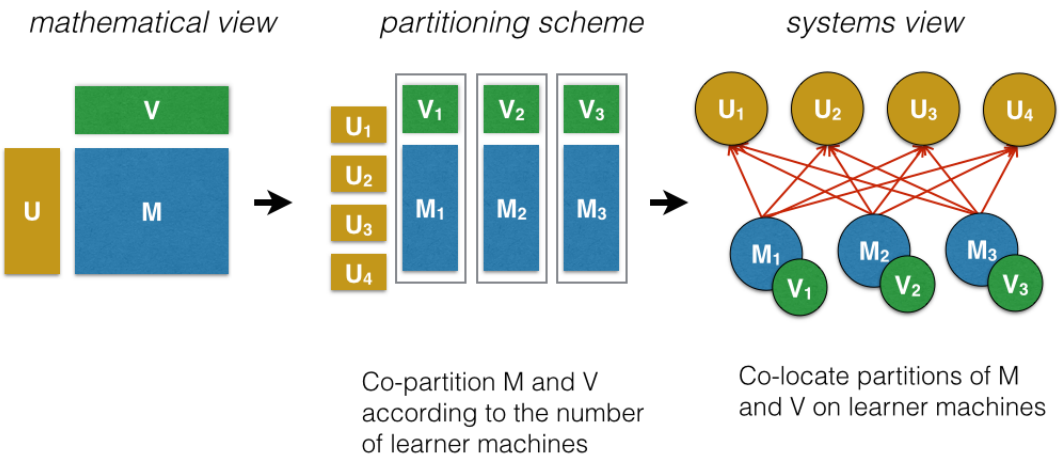

Figure 3: Partitioning scheme and system architecture.

4 Implementation

We implement Factorbird using twitter's existing technology stack and infrastructure. We leverage an existing memcached cluster for the parameter machines and implement the learner machines in Scala as a custom finagle application [18]. We use Hadoop's distributed filesystem (HDFS) as persistent datastore for reading inputs and storing the final results. Prepartitioning and statistics computation of the input is conducted via MapReduce using Scalding jobs. The learner machines are assigned resources via Apache Mesos.
The typical execution steps on a learner machine are as follows. First, the machine loads statistics about the input graph and its assigned partition from HDFS, e.g. the number of edges and vertices, the average edge weight and the degree distribution of the graph. These statistics are used for learning as well as efficiently allocating memory for the required datastructures. Next, the learner machine instantiates its local partition of V. Subsequently, the model learning phase starts: the learner machine reads the edges of its assigned partition of the input graph in a streaming fashion. For each edge (i, j), it reads the corresponding factor vector u_i from memcached and the factor vector v_j from its local partition of V. The reads from memcached are conducted in batches to increase throughput. Next, the learner machine updates the factor vectors using SGD and writes them back[1]. This process is repeated until a user-specified number of passes over the input graph has been conducted. Finally, the learner machines persist the learned matrices U and V in a partitioned manner in HDFS.
Factorbird makes extensive use of memory-efficient data structures for managing factor vectors, the

[1]Note that strictly speaking, we do not run vanilla Hogwild! on U, as memcached atomically updates the whole factor vector instead of individually updating the factors.

local partition of V and graph statistics such as the degree per vertex. The main objective of these data structures is to use as little memory as possible as well as to avoid object allocation and full garbage collection invocations in the JVM. A factor vector as well as a partition of a factor matrix are therefore internally represented by byte buffers and large primitive arrays, which are efficiently preallocated. For example, a learning machine determines the size of its local partition of V at startup time by reading the number of vertices assigned to its partition from the graph statistics. The partition of V is internally represented by a huge float array, into which the individual factor vectors (the columns of V) are packed. A mapping of vertex ids to offsets in this array is stored and update operations directly write to the underlying array. During the training phase, a learner machine directly reads the training edges from a compressed file in HDFS in a streaming fashion. If more than one pass through the training edges is necessary, the learner machine will on-the-fly create a copy of the HDFS file on local disk and switch to streaming edges from this local file for subsequent passes. Furthermore, some learning approaches require synthetically generated negative examples [24] (possibly taken into account with a lower confidence than observed positive examples). We therefore implement on-the-fly generation of such negative examples (with a configurable probability) and mix them into the original positive examples supplied on the learning machines.
The central abstraction for implementing the SGD updates of the learning algorithm within Factorbird is the *Learner* (c.f, Listing 1). It is the main interface for data scientists wanting to try new models and shields the programmer from the complexity of the distributed nature of the underlying learning process. The first important method that a programmer has to implement is *initialize* which Factorbird uses to randomly initialize of the factor vectors in U and V (c.f. line in Algorithm 1). The method *update* is where the SGD-based update of two factor vectors u_i and v_j is implemented. The system provides the strength $a(i, j)$ of the edge between vertex i and j, as well the vertex degrees n_i and n_j and the error weight $w(i, j)$ as additional arguments. A typical implementation of this method conducts the steps from line 4 to 9 in Algorithm 1.

```
trait Learner {
  def initialize(factors: FactorVector): Unit
  def update(u_i: FactorVector, v_j: FactorVector,
             a_ij: Float, n_i: Int, n_j: Int, w_ij: Float): Float
}
```

Listing 1: Learner abstraction.

5 Assessing Model Quality

The next aspect we focus on in Factorbird is the quality of the learned models. Ultimately, the models have to be evaluated using online experiments with real users, but during the development and batch training phase, we concentrate on a simple offline metric: the prediction quality on held-out data, measured by the root mean squared error (RMSE). Unfortunately, this prediction quality is heavily influenced by the choice of hyperparameters for our model, such as η (which controls rate of learning), λ (which controls regularization), number of factors k as well as the rate of decay of the learning rate in our *Learner* implementation. In order to find a well-working hyperparameter combination, we conduct a grid search in the hyperparameter space. We extend Factorbird to enable hold-out tests at scale. The Scalding job which prepares the input graph randomly splits the edges into training set, validation set and test set. Factorbird then learns a model on the training set, chooses the hyperparameter combination using the prediction quality on the validation set, and finally computes the RMSE on the test set.
However, conducting a single training run with Factorbird for each hyperparameter combination to inspect is tedious and takes a long time. We therefore we describe how to learn many models with different hyperparameters at once to speed up the hyperparameter search. Given that we aim to inspect c hyperparameter combinations for a factorization of rank k, we pack the c factor vectors into a large U of dimensionality $m \times c * k$ and a large V of dimensionality $c * k \times n$. We use a specialized learner implementation that is aware of the packing (e.g. it knows that the factors for the p-th model are contained in the $p * k$-th to $p * (k + 1)$-th entries of a factor vector) and learns c models at once. Figure 4 illustrates how the factor matrices for all possible combinations of two different learning rates η_1, η_2 and two different regularization constants λ_1, λ_2 would be packed into a large U and V.

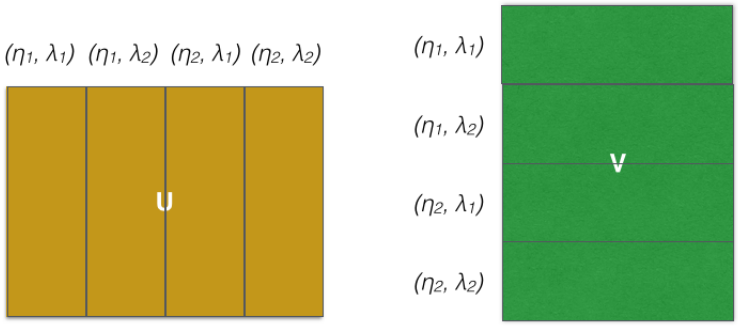

Figure 4: Packing many models into one for hyperparameter search.

In order to compute the prediction quality on held-out data, Factorbird asks the *Learner* implementation for *Predictor*s which the user has to provide. A *Predictor* predicts the strength of an unobserved edge (i, j) from the corresponding factor vectors u_i and v_j (c.f., Listing 2).

```
trait Predictor {
  def predict(u_i: FactorVector, v_j: FactorVector): Float
}
```

Listing 2: Predictor abstraction.

Optionally, the user can provide a *LossEstimator*, which estimates the current value of the loss function using samples of edges and factors (c.f., Listing 3). During training, Factorbird continually invokes the loss estimator and make the estimates inspectable via an external dashboard.

```
trait LossEstimator {
  def estimateRegularizationComponent(
    numRowsOfU: Int, sampleOfU: Iterator[FactorVector],
    numColumnsOfV: Int, sampleOfV: Iterator[FactorVector]): Double

  def estimateErrorComponent(numEdges: Long,
    sampleOfEdges: Iterator[Edge], partitionOfU: FactorMatrix,
    partitionOfV: FactorMatrix): Double
}
```

Listing 3: Abstraction for loss estimation.

6 Experiments

We run experiments on various subsets of 'RealGraph', a graph that models various interactions between twitter users [20]. The learner machines for our experiments are provisioned by Apache Mesos. In all our experiments, we factorize the binarized adjacency matrix of the graph subset. That means the, transformation function $a(i, j)$ returns 1 if user i interacted with user j and 0 otherwise (in the case of a synthetic negative example). We equally weight all prediction errors ($w(i, j) = 1$). In this work, we present only preliminary experiments, aimed at validating the correctness of our system and showing its capacity to handle twitter-scale graphs. There is still a large potential for improvements in accuracy and performance that we will tackle in future work. We run a first set of experiments on a small sample of the RealGraph, consisting of 100 million interactions between 440 thousand popular users. Additionally, we make Factorbird generate 500 million synthetic negative edges.

Benefits of increasing model complexity. In the first experiment, we show the positive effects on prediction quality of the individual parts of our chosen model. We randomly split the dataset into 80% training set, 10% validation set and 10% test set. We train models with increasing complex-

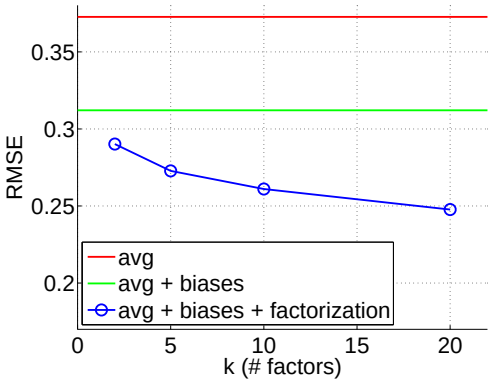

Figure 5: Prediction quality on held-out data with increasing model complexity.

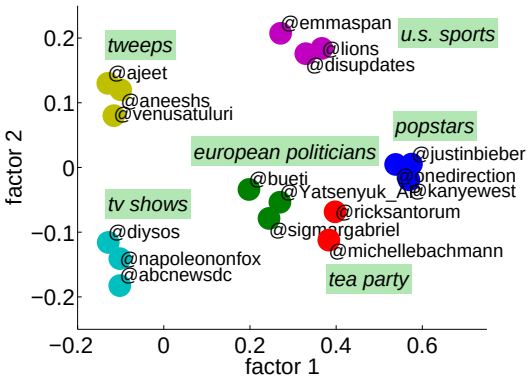

Figure 6: Plot of a selection of twitter users as positioned by a factorization with $k = 2$ of a sample of RealGraph.

ity and measure their prediction quality in terms of RMSE on the 10% held-out data in the test set (c.f. Figure 5). We start with a baseline that only uses the global bias (the average edge strength), followed by a more complex model that uses the global bias as well as vertex-specific bias terms. Finally, we train biased factorization models with $k \in \{2, 5, 10, 20\}$. We choose the hyperparameters using the validation set. The outcome confirms that an increase in the complexity of our model results in an increase in prediction quality. The global bias baseline provides an RMSE of 0.3727, adding the vertex-specific bias terms reduces the error to 0.3121 and incorporating the factors gives additional improvements, reducing the error to 0.2477 for $k = 20$.

Visual inspection. Next, we plot a selection of factor vectors from the V matrix of a biased factorization with $k = 2$. Two users i and j will be close in the resulting low-dimensional space if their follower vectors $m_{\bullet i}$ and $m_{\bullet j}$ are roughly linear combinations of each other. Due to homophily, we expect twitter users sharing a common interest or characteristic to have a large overlap in followers and therefore to be much closer in the resulting space. Figure 6 indicates that factorizations produced by Factorbird have this property. We see several clusters of twitter accounts of similar type, e.g. a cluster of european politicians, containing Sigmar Gabriel (`@sigmargabriel`) and Reinhard Buetikofer (`@bueti`) from Germany as well as Arseniy Yatsenyuk (`@Yatsenuk_AP`) from Ukraine. Another cluster consists of popstars such as Kanye West (`@kanyewest`), Justin Bieber (`@justinbieber`) and the boy-band *One Direction* (`@onedirection`). A third one related to U.S. sports contains Emma Span (`@emmaspan`), an editor for baseball at Sports Illustrated, Detroit's football team (`@lions`) and an account of the Daytona International Speedway (`@disupdates`).

Scale-out. For the scale-out experiments we build a matrix based on a large subset of Twitter's RealGraph [20] consisting of more than 38.5 billion non-zeros (half of which are synthetic negative examples). The dimensions of the matrix are 229 million $\times$ 195 million. To the best of our knowledge, this is the largest dataset on which collaborative filtering experiments have been published so far. Often, the Netflix prize dataset [3] is used for scale-out tests [11]. Our dataset is more than two orders of magnitude larger, having approximately 470 times more rows, 11,000 times more columns and 385 times more datapoints. We run Factorbird using 50 instances for the learner machines, provisioned in a Mesos cluster, with 16 cores and 16 gb of memory each. We leverage a shared memcached cluster for our parameter machines, with a guaranteed quota of 5 million commands per second. We run two passes of hyperparameter search on 80% of the data for 16 different learners with k = 5. The rank of the factor matrices U and V is $(5 + 1) * 16 = 96$, which means we train a model with approximately $(229M + 195M) * 96 \approx 40B$ parameters. A single SGD pass through the training set finishes in about 2.5 hours [2]. In this experiment, Factorbird issues more than 4.2 million commands per second to memcached on average and updates around 400M parameters per second.

[2]Except for a few stragglers, which Mesos provisioned on machines with low network bandwidth (Unfortunately, Mesos cannot guarantee a defined network bandwidth at the moment)

7 Related Work

SGD- and ALS-based matrix factorization techniques for recommender systems have been extensively studied in the context of the Netflix Prize [2, 9]. These techniques have been extended to work on implicit feedback data [8] and to optimize metrics different from RMSE [10, 16].
A large body of work has been conducted with respect to parallelizing and distributing matrix factorization. This includes work on the scalability of the algorithms itself, e.g. by introducing biased sampling for SGD to avoid conflicting updates during concurrent execution [7, 5, 15] or by proving convergence under a minor amount of update conflicts [14]. Furthermore, distributed implementations have been proposed and evaluated on MapReduce-based [4], graph-parallel [11, 12] and specialized systems [13].

8 Future Work

In future work, we aim to extend Factorbird to a streaming scenario. We plan to bootstrap Factorbird with a factorization that was trained offline and update this factorization model online from a stream of incoming real-time interactions (e.g., user follows). Furthermore, we would like to experiment with factorizing multiple matrices at once, in order to incorporate different types of interactions in a single factorization model. Another technical issue to work on in the future is fault tolerance. A possible approach to recovery in case of failures could be to restart learner and parameter machines from asynchronously written checkpoints of the partitions of U and V which they hold [19]. Moreover, we plan to investigate ways to reduce the amount of network traffic caused by Factorbird, e.g. by compression, factor vector caching or via a biased sampling of the edges to allow us to use retrieved factor vectors for more than a single update. We will potentially replace memcached with a custom application to be able to achieve higher throughput and conduct true Hogwild-style updates on the parameter machines. Moreover, this would allow us to run aggregations on the parameter machines. Additionally, we would like to implement dynamic load adaption in Factorbird to mitigate the negative effects of stragglers on the overall runtime. We aim to factorize the whole twitter interaction graph with Factorbird in upcoming work.

References

[1] Amatriain, X. (2012) Building industrial-scale real-world recommender systems. In *ACM Conference on Recommender Systems*, http://goo.gl/cqKAi

[2] Koren, Y., Bell, R. & Volinsky C. (2009) Matrix factorization techniques for recommender systems. In *Computer 42.8*, pp. 30–37.

[3] Bennett, J. & Lanning, S. (2007) The Netflix Prize. In *KDD Cup and Workshop*, p. 35.

[4] Schelter, S., Boden, C., Schenck, M., Alexandrov, A. & Markl, V. (2013). Distributed matrix factorization with mapreduce using a series of broadcast-joins. In *ACM Conference on Recommender Systems*, pp. 281-284.

[5] Teflioudi, C., Makari, F. & Gemulla, R. (2012) Distributed Matrix Completion. In *IEEE International Conference on Data Mining*, pp. 655-664.

[6] Das, A. S., Datar, M., Garg, A. & Rajaram, S. (2007). Google news personalization: scalable online collaborative filtering. In *International Conference on the World Wide Web*, pp. 271-280.

[7] Gemulla, R., Nijkamp, E., Haas, P. J. & Sismanis, Y. (2011). Large-scale matrix factorization with distributed stochastic gradient descent. In *ACM Conference on Knowledge Discovery and Data mining*, pp. 69-77. top of the rock new york

[8] Hu, Y., Koren, Y. & Volinsky, C. (2008). Collaborative filtering for implicit feedback datasets. In *IEEE International Conference on Data Mining*, pp. 263-272.

[9] Zhou, Y., Wilkinson, D., Schreiber, R. & Pan, R. (2008). Large-scale parallel collaborative filtering for the netflix prize. In *Algorithmic Aspects in Information and Management*, pp. 337-348.

[10] Rendle, S., Freudenthaler, C., Gantner, Z. & Schmidt-Thieme, L. (2009). BPR: Bayesian personalized ranking from implicit feedback. In *Conference on Uncertainty in Artificial Intelligence*, pp. 452-461.

[11] Low, Y., Bickson, D., Gonzalez, J., Guestrin, C., Kyrola, A. & Hellerstein, J. M. (2012). Distributed GraphLab: a framework for machine learning and data mining in the cloud. In *Proceedings of the VLDB Endowment*, 5(8), pp. 716-727.

[12] Xin, R. S., Crankshaw, D., Dave, A., Gonzalez, J. E., Franklin, M. J. & Stoica, I. (2014). GraphX: Unifying Data-Parallel and Graph-Parallel Analytics. *arXiv preprint arXiv:1402.2394*.

[13] Yun, H., Yu, H. F., Hsieh, C. J., Vishwanathan, S. V. N. & Dhillon, I. (2013). NOMAD: Non-locking, stOchastic Multi-machine algorithm for Asynchronous and Decentralized matrix completion. *Proceedings of the VLDB Endowment*, 7(11).

[14] Recht, B., Re, C., Wright, S. & Niu, F. (2011). Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, pp. 693-701.

[15] Recht, B. & Re, C. (2013). Parallel stochastic gradient algorithms for large-scale matrix completion. In *Mathematical Programming Computation*, 5(2), pp. 201-226.

[16] Weimer, M., Karatzoglou, A., Le, Q. V. & Smola, A. (2007). Maximum Margin Matrix Factorization for Collaborative Ranking. In *Advances in Neural Information Processing Systems*.

[17] Li, M., Zhou, L., Yang, Z., Li, A., Xia, F., Andersen, D. G. & Smola, A. (2014) Parameter Server for Distributed Machine Learning.

[18] Eriksen, M. (2013). Your server as a function. In *Workshop on Programming Languages and Operating Systems*, p. 5.

[19] Schelter, S., Ewen, S., Tzoumas, K. & Markl, V. (2013). "All Roads Lead to Rome:" Optimistic Recovery for Distributed Iterative Data Processing. In *ACM International Conference on Information & Knowledge Management*, pp. 1919-1928.

[20] Kamath, K., Sharma, A., Wang, D. & Yin, Z. (2014) RealGraph: User Interaction Prediction at Twitter. In *User Engagement Optimization Workshop @ KDD*

[21] Myers, S. A., Sharma, A., Gupta, P. & Lin, J. (2014). Information network or social network?: the structure of the twitter follow graph. In *International Conference on the World Wide Web*, pp. 493-498.

[22] Hastie, T., Mazumder, R., Lee, J. D. Lee & Zadeh, R. (2014). Matrix Completion and Low-Rank SVD via Fast Alternating Least Squares. In *arXiv*

[23] Udell, M., Horn, C., Zadeh, R. & Boyd, S. (2014). Generalized Low Rank Models. In *arXiv*

[24] Pan, R., Zhou, Y., Cao, B., Liu, N. N., Lukose, R., Scholz, M. & Yang, Q. (2008). One-class collaborative filtering. In *IEEE International Conference on Data Mining*, pp. 502-511.